

CLAIMS

1 1. A method for evaluating the locality of references for a
2 layout of a computer program, the computer program to be executed by a
3 computer system with a page architecture, each page having a multiplicity of
4 memory locations, the method comprising:

5 for each of a plurality of memory locations of a page, calculating a
6 metric value indicating a working set size of the layout when the layout is
7 positioned to start at that memory location; and

8 combining the calculated metric values as an indication of the
9 locality of references of the layout of the computer program.

1 2. The method of claim 1 wherein the plurality of memory
2 locations include each memory location of a page.

1 3. The method of claim 1 wherein the plurality of memory
2 locations include less than all the memory locations on a page.

1 4. The method of claim 1 wherein the plurality of memory
2 locations are separated by relatively prime number of memory locations.

1 5. The method of claim 1 wherein the plurality of memory
2 locations includes less than all the memory locations on a page and are
3 approximately evenly distributed throughout the page.

1 6. The method of claim 1 wherein the plurality of memory
2 locations includes less than all the memory locations on a page, are
3 approximately evenly distributed throughout the page, and are separated by
4 relatively prime number of memory locations.

1 7. The method of claim 1 wherein the combining of the
2 calculated metric values includes averaging the calculated metric values.

1 8. A method in a computer system for determining a number of
2 different layouts of a program image that should be evaluated when selecting a
3 layout as a starting layout for a working set optimization process, the method
4 comprising:

5 collecting data on results of using the working set optimization
6 process to improve the working set of various layouts of various program images;
7 and

8 determining from the collected number of layouts whose
9 computational expense in generating and evaluating the layouts is worth an
10 anticipated reduction in the working set when the program image is optimized by
11 the working set optimization process.

1 9. The method of claim 8 wherein the collecting of data
2 includes calculating a marginal density of the working set size of the optimized
3 layout that is produced from one of the various layouts with the smallest locality
4 of reference.

1 10. The method of claim 8 wherein the collecting of data
2 includes calculating a mean marginal density that is

$$\mu = \int_{-\infty}^{\infty} y g(y) dy$$

6 where

7

$$8 \quad g(y) = N \int_{-\infty}^{\infty} f(x, y) \left(\int_{-\infty}^{\infty} \int_x^{\infty} f(t, u) dt du \right)^{N-1} dx$$

9

10 where N is the number of layouts and f is a joint probability density function that
11 relates a metric for each unoptimized layout to a metric for each corresponding
12 optimized layout.

1

11. The method of claim 10 wherein the joint probability
2 density function is

3

$$4 \quad f(x, y) = \frac{1}{2\pi\sqrt{1-\rho^2}} e^{-\frac{(x^2-2\rho xy+y^2)}{2(1-\rho^2)}}$$

5

6 and where ρ is a normal correlation coefficient between a locality of reference
7 metric for the unoptimized layouts and a working set size metric for the
8 optimized layouts.

1

12. The method of claim 10 wherein the anticipated reduction is

2

3

$$\mu\sigma$$

4

5 where σ is a standard deviation of a working set size metric for the optimized
6 layouts.

1

13. The method of claim 10 wherein the probability density
2 function is a standard bivariate normal distribution.

1

14. The method of claim 8 wherein the determining of the
2 number of layouts includes determining the computational expense of generating

3 and evaluating successively larger number of layouts until the computation
4 expense is not worthwhile.

1 15. A method for estimating a rate of change per time period in
2 results of a system, the system producing a result for each of a plurality of steps
3 performed by the system, the steps occurring at irregular time intervals, the
4 method comprising:

5 estimating the change in results at each of the plurality of steps;

6 estimating the time for each step; and

7 combining the estimated change in results per step with the
8 estimated time per step to estimate the rate of change per time period in the
9 results of the system.

1 16. The method of claim 15 wherein the estimating of the
2 change in results includes filtering the actual change in results from one step to
3 the next step.

1 17. The method of claim 15 wherein the estimating of the time
2 for each step includes evaluating a formula that estimates the time per step based
3 on estimated numbers of times that various sub-steps are performed per step.

1 18. The method of claim 17 wherein the estimated numbers of
2 times that various sub-steps are performed are based on evaluating actual results
3 of the system.

1 19. The method of claim 15 wherein the estimating of the
2 change in results includes filtering the actual change in results from one step to
3 the next step and wherein the estimating of the time for each step includes

4 evaluating a formula that estimates the time per step based on number of time
5 various sub-steps are performed.

1 20. The method of claim 15 wherein the system is a computer
2 program layout optimization system.

1 21. The method of claim 15 including terminating the system
2 when the estimated rate of change per time period is outside of a threshold rate of
3 change.

1 22. The method of claim 15 wherein the estimating of the
2 change in the results includes filtering the actual change in results from one step
3 to the next with a filter that is generated from analysis of results produced by the
4 system.

1 23. The method of claim 22 wherein the filter is generated by
2 using a frequency-domain analysis of an actual rate of change per step and
3 defined rate of change per step, the defined rate of change being calculated for
4 each step based on knowledge of results of subsequent steps.

1 24. The method of claim 22 wherein the filter is generated by
2 using a time-domain analysis of an actual rate of change per step and defined rate
3 of change per step, the defined rate of change being calculated for each step
4 based on knowledge of results of subsequent steps.

1 25. A method in a computer system for estimating the rate of
2 improvement in the working set for a plurality of incrementally improved layouts
3 of a computer program, the method comprising:

4 estimating the change in working set size from one incrementally
5 improved layout to the next incrementally improved layout;
6 estimating the time needed to incrementally improve the layout;
7 and
8 combining the estimated change in working set size with the
9 estimated time needed to incrementally improve the working set for that layout to
10 estimate the rate of improvement.

1 26. The method of claim 25 wherein the estimating of the
2 change in working set size includes filtering the actual change in working set size
3 from one incremental improvement to the next.

1 27. The method of claim 25 wherein the estimating of the time
2 to incrementally improve the layout includes evaluating a formula that estimates
3 the time based on estimated numbers of times that various sub-steps are
4 performed when incrementally improving the layout.

1 28. The method of claim 27 wherein the estimated numbers of
2 times that various sub-steps are performed are based on evaluating actual results
3 of the incremental improvement.

1 29. The method of claim 25 wherein the estimating of the
2 change includes filtering the actual change in working set size from one
3 incremental improvement to the next and wherein the estimating of the time to
4 incrementally improve the layout includes evaluating a formula that estimates the
5 time based on number of times various sub-steps are performed when
6 incrementally improving the layout.

1 30. The method of claim 25 wherein the estimating of the
2 change in working set size filters an actual change in working set size and the
3 estimating of time computes the time based on number of times the layout has
4 been incrementally improved.

1 31. The method of claim 25 including terminating the system
2 when the estimated rate of change per time period is outside of a threshold rate of
3 change.

1 32. The method of claim 25 wherein the estimating of the
2 change in the working set size includes filtering the actual change in working set
3 size from one incremental improvement to the next with a filter that is generated
4 from analysis of working set size produced during the incremental improvement
5 of other layouts.

1 33. The method of claim 32 wherein the filter is generated by
2 using a frequency-domain analysis of an actual rate of change per incremental
3 improvement and defined rate of change per incremental improvement, the
4 defined rate of change being calculated for each incremental improvement based
5 on knowledge of working set size of subsequent incremental improvements.

1 34. The method of claim 32 wherein the filter is generated by
2 using a time-domain analysis of an actual rate of change per incremental
3 improvement and defined rate of change per incremental improvement, the
4 defined rate of change being calculated for each incremental improvement based
5 on knowledge of working set size of subsequent incremental improvements.

1 35. A method in a computer system for estimating a rate of
2 change of results of a function over an indeterminate interval, the method
3 comprising:
4 receiving a plurality of results of the function;
5 calculating a negative slope of a convex hull of the results of the
6 function; and
7 using the calculated negative slope as an estimate of the rate of
8 change of results of the function.

1 36. The method of claim 35 wherein the results of the function
2 are generated by calculating a working set size of a program image as the layout
3 is incrementally improved.

1 37. The method of claim 35 including calculating a running
2 minimum of the results of the function wherein the calculating of the negative
3 slope of a convex hull does so for the calculated running minimum.

1 38. The method of claim 37 wherein the results of the function
2 are generated by calculating a working set size of a program image as the layout
3 is incrementally improved.

1 39. A method in a computer system for identifying coefficients
2 for a filter, the filter for filtering results of a function, the method comprising:
3 collecting sample input values to the filter;
4 identifying desired output values from the filter for the collected
5 sample input values;
6 generating a power spectrum of the collected sample input values;

7 generating a power spectrum of the identified desired output
8 values;
9 calculating a difference between the generated power spectra; and
10 identifying coefficients that minimize differences between the
11 calculated differences.

1 40. The method of claim 39 wherein the calculated difference is
2 offset so that the calculated difference at a frequency of zero is zero.

1 41. The method of claim 39 wherein the results of the function
2 are generated by calculating a working set size of a program image as the
3 working set size is incrementally improved.

1 42. The method of claim 39 wherein a Levenberg-Marquardt
2 algorithm is used to identify the coefficients.

1 43. A method in a computer system for identifying coefficients
2 for a finite impulse response filter, the method comprising:
3 collecting sample input values for a function;
4 identifying desired output values for the filter for the collected
5 sample input values;
6 approximating the output values from the input values using a
7 linear fitting technique; and
8 setting the coefficients to values obtained from the linear-fitting
9 technique.

1 44. The method of claim 43 wherein the results of the function
2 are generated by calculating a working set size of a program image as the
3 working set size is incrementally improved.

1 45. The method of claim 43 wherein the linear-fitting technique
2 is for a certain-order finite impulse response filter and including determining an
3 error between the approximated output values and the desired output values and
4 when the difference between the determined error and input for the next higher-
5 order finite impulse response filter is not significant using the certain-order finite
6 impulse response filter.

1 46. The method of claim 45 wherein the significance of the
2 difference between the determined error and input for the next higher-order finite
3 impulse response filter is determined by correlating the determined error with the
4 additional input for the next higher-order finite impulse response filter.

1 47. A method for converting a finite impulse response filter into
2 an infinite impulse response filter so that a desired tradeoff between confidence
3 and mean lag is established, the method comprising:
4 for each of a plurality of sets of autoregressive coefficients,
5 calculating results of the infinite impulse response filter with
6 the set of autoregressive coefficients, and
7 generating a value that rates the confidence and mean lag for
8 the calculated results; and
9 selecting the set of autoregressive coefficients with the generated
10 value with the highest rating of confidence and mean lag for the infinite impulse
11 response filter.

1 48. The method of claim 47 wherein the filters are for
2 determining a rate of improvement in a working set optimization process.

1 49. The method of claim 47 wherein the plurality of sets are
2 generated using Brent's method.

1 50. The method of claim 47 wherein the plurality of sets are
2 generated using Powell's method.

1 51. A computer-readable medium containing instructions for
2 causing a computer system to evaluate locality of references for a layout of a
3 computer program, the computer program to be executed by a computer system
4 with a page architecture, each page having memory locations, by:

5 for each of a plurality of selected memory locations of a page,
6 estimating a working set size of the layout when the layout is positioned to start
7 at that memory location; and

8 combining the estimated working set size as an indication of the
9 locality of references of the layout of the computer program.

1 52. The computer-readable medium of claim 51 wherein the
2 plurality of selected memory locations include each memory location of a page.

1 53. The computer-readable medium of claim 51 wherein the
2 plurality of selected memory locations include less than all the memory locations
3 on a page.

1 54. The computer-readable medium of claim 51 wherein the
2 plurality of selected memory locations are separated by relatively prime number
3 of memory locations.

1 55. The computer-readable medium of claim 51 wherein the
2 plurality of selected memory locations includes less than all the memory

3 locations on a page and are approximately evenly distributed throughout the
4 page.

1 56. The computer-readable medium of claim 51 wherein the
2 plurality of selected memory locations includes less than all the memory
3 locations on a page, are approximately evenly distributed throughout the page,
4 and are separated by relatively prime number of memory locations.

1 57. The computer-readable medium of claim 51 wherein the
2 combining of the estimated working set size includes averaging the estimated
3 working set sizes.

1 58. A computer system for estimating the rate of improvement
2 in the working set size for a plurality of layouts of a computer program, the
3 layouts resulting from a layout optimization process, comprising:
4 a first estimating component that estimates the change in working
5 set size from one improved layout to the next improved layout;
6 a second estimating component that estimates the time needed to
7 improve the layout; and
8 a combining component that combines the estimated change in
9 working set size with the estimated time needed to improve the working set for
10 that layout to estimate the rate of improvement.

1 59. The system of claim 58 wherein the estimating of the change
2 in working set size includes filtering the actual change in working set size from
3 one layout to the next.

1 60. The system of claim 58 wherein the estimating of the time to
2 improve the layout includes evaluating a formula that estimates the time based on

3 estimated numbers of times that various sub-steps are performed when improving
4 the layout.

1 61. The system of claim 60 wherein the estimated numbers of
2 times that various sub-steps are performed are based on evaluating actual results
3 of the improvement.

1 62. The system of claim 58 wherein the estimating of the change
2 includes filtering the actual change in working set size from one layout to the
3 next and wherein the estimating of the time to improve the layout includes
4 evaluating a formula that estimates the time based on number of times various
5 sub-steps are performed when improving the layout.

1 63. The system of claim 58 wherein the estimating of the change
2 in working set size filters an actual change in working set size and the estimating
3 of time computes the time based on number of times the layout has been
4 improved.

1 64. The system of claim 58 including terminating the system
2 when the estimated rate of change per time period is outside of a threshold rate of
3 change.

1 65. The system of claim 58 wherein the estimating of the change
2 in the working set size includes filtering the actual change in working set size
3 from one layout to the next with a filter that is generated from analysis of
4 working set size produced during the improvement of other layouts.

1 66. The system of claim 65 wherein the filter is generated by
2 using a frequency-domain analysis of an actual rate of change per improvement

3 and defined rate of change per improvement, the defined rate of change being
4 calculated for each improvement based on knowledge of working set size of
5 subsequent improvements.

1 67. The system of claim 65 wherein the filter is generated by
2 using a time-domain analysis of an actual rate of change per improvement and
3 defined rate of change per improvement, the defined rate of change being
4 calculated for each improvement based on knowledge of working set size of
5 subsequent improvements.

1 68. A computer-readable medium containing instructions for
2 causing a computer system to convert a finite impulse response filter into an
3 infinite impulse response filter so that a desired tradeoff between confidence and
4 mean lag is established, by:

5 for each of a plurality of sets of autoregressive coefficients, for the
6 infinite impulse response filter calculating results of the infinite impulse response
7 filter with the set of autoregressive coefficients; and

8 selecting the set of autoregressive coefficients with a highest rating
9 of the trade off between confidence and mean lag for the infinite impulse
10 response filter.

1 69. The computer-readable medium of claim 68 wherein the
2 filters are for determining a rate of improvement in a working set optimization
3 process.

1 70. The computer-readable medium of claim 68 wherein the
2 plurality of sets are generated using Brent's method.

- 1 71. The computer-readable medium of claim 68 wherein the
- 2 plurality of sets are generated using Powell's method.